A near-optimal sampling strategy for sparse recovery of polynomial chaos expansions

Negin Alemazkoor, Hadi Meidani

Department of Civil and Environmental Engineering, University of Illinois at Urbana-Champaign, Urbana, Illinois, USA.

arXiv:1702.07830v1 [stat.CO] 25 Feb 2017

Abstract

Compressive sampling has become a widely used approach to construct polynomial chaos surrogates when the number of available simulation samples is limited. Originally, these expensive simulation samples would be obtained at random locations in the parameter space. It was later shown that the choice of sample locations could significantly impact the accuracy of resulting surrogates. This motivated new sampling strategies or design-of-experiment approaches, such as coherence-optimal sampling, which aim at improving the coherence property. In this paper, we propose a sampling strategy that can identify near-optimal sample locations that lead to improvement in local-coherence property and also enhancement of cross-correlation properties of measurement matrices. We provide theoretical motivations for the proposed sampling strategy along with several numerical examples that show that our near-optimal sampling strategy produces substantially more accurate results, compared to other sampling strategies.

1. Introduction

In order to facilitate stochastic computation in analysis and design of complex systems, analytical surrogates that approximate and replace full-scale simulation models have been increasingly studied [1, 2, 3]. One of the most widely adopted surrogates is the polynomial chaos expansion (PCE), which approximates the quantity of interest (QoI) by a spectral representation using polynomial functions of random parameters [4, 5, 6, 7]. In estimating these spectral surrogates, non-intrusive stochastic techniques, based on either spectral projection or linear regression, are widely used especially because they don't require modifying deterministic solvers or legacy codes, which is an otherwise cumbersome task [8]. These non-intrusive techniques are still the subject of ongoing research as the number of required samples for accurate surrogate estimation rapidly grows with the number of random parameters, even when efficient techniques such as sparse grid are used [9, 10, 11, 12].

More recently, researchers have developed techniques, based on compressive sampling (CS), that are particularly advantageous when surrogate expansions are expected to be sparse, i.e. the QoI can be accurately represented with a few polynomial chaos (PC) bases. Compressive sampling was first introduced in the field of signal processing to recover sparse signals using a number of samples significantly smaller that the conventionally used Shannon-Nyquist sampling rate [15, 16, 17]. Motivated by the fact that the solution of many high dimensional problems of interest, such as high dimensional PDEs, can be represented by sparse, or at least approximately sparse, PCEs, CS was proposed in [18, 19, 20] to estimate PC coefficients in underdetermined cases. As CS theorems suggest, the success of sparse estimation of PCE depends not only upon the sparsity of the solution of stochastic system, but also on incoherence property of the Vondermonde-like measurement matrix, formed by evaluations of orthogonal polynomials at sample locations [15], as will be elaborated later.

Several efforts have been made in order to improve the two mentioned conditions for successful recovery. For instance in [21], for Hermite expansions with Gaussian input variables, the original inputs are rotated such that a few of the new coordinates, i.e. linear combinations of original inputs, have significant impact on QoI, thereby increasing the sparsity of solution and, in turn, the accuracy of recovery. The section condition, i.e. incoherence of the measurement matrix, can be poor especially when trial expansions are high-order and/or high-dimensional. To remedy this, the iterative approaches in [22, 23] can be used to optimally include only the "important" bases into the trial expansion and its associated measurement matrix. Focusing on this second condition, another class of methods have proposed sampling strategies that produce less coherent measurement matrices [24, 25, 26]. Among these approaches, the sampling strategy proposed in [26] was designed to be optimal in achieving the lowest local-coherence.

In this work, we introduce a near-optimal sampling strategy by further improving the local-coherence-based sampling of [26] and filtering sample locations based on cross-correlation properties of the resulting measurement matrix. Specifically, we establish quantitative measures to capture these cross-correlations properties between measurement matrix columns, and use these measures as the criteria for near-optimal identification of sample locations. It will be demonstrated that a sampling strategy that seeks to optimize these measures will lead to CS results that on average outperforms all other CS sampling strategies. This paper is organized as follows. Section 2 presents general concepts in compressive sampling and its theoretical background. In Section 3, we introduce our sampling algorithm along with relevant theoretical supports. Finally, Section 4 includes numerical examples and discussions about the advantages of the proposed approach.

2. Setup and background

2.1. Polynomial chaos expansion

PC expansion is a spectral technique to approximate stochastic models using orthogonal polynomials of random inputs. Formally put, let the d random inputs be denoted by the d-dimensional random vector $\boldsymbol{\Xi} = (\Xi_1, ..., \Xi_d)$ with mutually independent components, a joint density function denoted by $\rho(\boldsymbol{\xi})$, and finite moments in the probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where Ω is the sample space with its corresponding σ-algebra, $\mathcal{F}$, and probability measure, $\mathbb{P}$. Also, let $\rho_i(\xi_i)$ be the probability density function for Ξ_i, and Ω_i be its corresponding sample space such that

$$\rho(\boldsymbol{\xi}) = \prod_{i=1}^{d} \rho_i(\xi_i), \tag{1}$$

$$\Omega = \prod_{i=1}^{d} \Omega_i. \tag{2}$$

Then, one can introduce the set of one dimensional orthonormal polynomial $\{\psi_{\alpha,i}\}_{\alpha\in\mathbb{N}_0}$ in Ξ_i that satisfies

$$\int_{\Omega_i} \psi_{n,i}(\xi_i)\psi_{m,i}(\xi_i)\rho_i(\xi_i)\,\mathrm{d}\xi_i = \delta_{mn}, \quad m, n \in \mathbb{N}_0. \tag{3}$$

where $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$, and δ_{mn} is the delta function. Therefore, the density function of Ξ_i, $\rho_i(\xi_i)$, determines the type of polynomial. For example, Gaussian and uniform probability distributions enforce Hermite and Legendre polynomials, respectively. The d-dimensional orthonormal polynomials are then derived from the multiplication of one dimensional polynomials in all dimensions. For example,

$$\psi_{\boldsymbol{m}}(\boldsymbol{\xi}) = \psi_{m_{1,1}}(\xi_1)...\psi_{m_{d,d}}(\xi_d), \quad \boldsymbol{m} = (m_1, ..., m_d). \tag{4}$$

Consequently, we have

$$\int_{\Omega} \psi_{\boldsymbol{n}}(\boldsymbol{\xi})\psi_{\boldsymbol{m}}(\boldsymbol{\xi})\rho(\boldsymbol{\xi})\,\mathrm{d}\boldsymbol{\xi} = \delta_{\boldsymbol{mn}}, \quad \boldsymbol{m}, \boldsymbol{n} \in \mathbb{N}_0^d. \tag{5}$$

Using this construction, any function $u(\boldsymbol{\Xi}) : \Omega \to \mathbb{R}$ that is square-integrable can be represented as

$$u(\boldsymbol{\Xi}) = \sum_{\boldsymbol{\alpha}\in\mathbb{N}_0^d} c_{\boldsymbol{\alpha}}\psi_{\boldsymbol{\alpha}}(\boldsymbol{\Xi}), \tag{6}$$

where $\{\psi_{\boldsymbol{\alpha}}\}_{\boldsymbol{\alpha}\in\mathbb{N}_0^d}$ is the set of orthonormal basis functions satisfying Equation (5). However, for computation's sake, $u(\boldsymbol{\Xi})$ is approximated by a finite order truncation of PC expansion given by

$$u_k(\boldsymbol{\Xi}) := \sum_{\boldsymbol{\alpha}\in\Lambda_{d,k}} c_{\boldsymbol{\alpha}}\psi_{\boldsymbol{\alpha}}(\boldsymbol{\Xi}), \tag{7}$$

where k is the total order of the polynomial expansion and $\Lambda_{d,k}$ is the set of multi-indices defined as

$$\Lambda_{d,k} := \{\boldsymbol{\alpha} \in \mathbb{N}_0^d : \|\boldsymbol{\alpha}\|_1 \leq k\}. \tag{8}$$

The cardinality of $\Lambda_{d,k}$, i.e. the number of expansion terms, here denoted by K, is a function of d and k according to

$$K := |\Lambda_{d,k}| = \frac{(k+d)!}{k!d!}. \tag{9}$$

Given this setting, $u_k(\boldsymbol{\Xi})$ approximates $u(\boldsymbol{\Xi})$ in a proper sense and is referred to as the k-th degree PC approximation of $u(\boldsymbol{\Xi})$. Each of the K coefficients involved in the definition of u_k can be exactly calculated by projecting $u(\boldsymbol{\Xi})$ onto the associated basis function:

$$c_{\boldsymbol{\alpha}^j} = \frac{1}{\gamma_{\boldsymbol{\alpha}^j}} \int_\Omega u(\boldsymbol{\xi}) \psi_{\boldsymbol{\alpha}^j}(\boldsymbol{\xi}) \rho(\boldsymbol{\xi}) \, \mathrm{d}\boldsymbol{\xi}, \quad \boldsymbol{\alpha}^j \in \Lambda_{d,k} = \{\boldsymbol{\alpha}^1, \cdots, \boldsymbol{\alpha}^K\}. \tag{10}$$

Numerical approaches such as cubature rule or sparse grid are usually used to approximate the above integral as its exact evaluation might be very cumbersome or not even possible [8, 27]. However, as the dimensionality of the problems increases, the number of samples required by these numerical techniques for an accurate integral evaluation increases exponentially. Supplying such a large sample size is prohibitively costly, especially if these samples are drawn from expensive high-fidelity simulations or costly experiments. As a result, efforts have been made to develop adaptive approaches that reduce the number of required samples by placing fewer samples on dimensions that do not impact QoI significantly [13, 14]. Another widely used approach to estimate PC coefficients is linear regression. In order to use linear regression, the number of samples must be equal or greater than the number of PC expansion terms. The well accepted oversampling rate is around 1.5 to 3 times the number of unknown coefficients. Recently, a quasi-optimal sampling approach was introduced in [28] that results in accurate regression with $\mathcal{O}(1)$ oversampling.

As briefly mentioned in Section 1, when the QoI is sparse with respect to the PC bases, compressive sampling approach have proved effective in estimating expansion coefficients using a number of samples that is significantly smaller than the number of PC terms. This section follows with the review on the basics of sparse PC recovery using compressive sampling.

2.2. Stochastic collocation using compressive sampling

Compressive sampling was first introduced in the field of signal processing, where conventionally the number of samples required to recover a signal was determined by Shannon-Nyquist sampling rate [15]. Compressive sampling allows for successful signal recovery using significantly fewer samples when the signal of interest is sparse. Therefore, it has been extensively applied in cases where the number of available samples is limited [29, 30, 31, 32]. Due to this very advantage, compressive sampling has recently gained substantial attention in UQ, more specifically in stochastic collocation, where a sample set is used to build an analytical surrogate model (typically in the form of a PC expansion) [18, 19, 26, 33, 34]. In what follows, a formal brief background on the use of compressive sampling for PCE estimation is provided.

The objective of PCE estimation is to calculate the vector of unknown expansion coefficients $\boldsymbol{c} = (c_{\boldsymbol{\alpha}^1}, .., c_{\boldsymbol{\alpha}^K})^T$ in Equation (7), given M samples of the QoI, denoted by the data vector $\boldsymbol{u} = (u(\boldsymbol{\xi}^{(1)}), ..., u(\boldsymbol{\xi}^{(M)}))^T$. These sampled outputs are evaluated at the M realizations, $\{\boldsymbol{\xi}^{(i)}\}_{i=1}^M$, of model input $\boldsymbol{\Xi}$. Requiring $u_k(\boldsymbol{\Xi})$ to approximate $u(\boldsymbol{\Xi})$ results in the following system of equation,

$$\boldsymbol{\Psi}\boldsymbol{c} = \boldsymbol{u}, \tag{11}$$

where $\boldsymbol{\Psi}$ is the measurement matrix, constructed according to

$$\boldsymbol{\Psi} = [\psi_{ij}], \quad \psi_{ij} = \psi_{\boldsymbol{\alpha}^j}(\boldsymbol{\xi}^{(i)}), \quad 1 \leqslant i \leqslant M, \quad 1 \leqslant j \leqslant K. \tag{12}$$

We are interested in the underdetermined case where $M \leqslant K$. Under such condition, there may exist infinity many solutions for $\boldsymbol{c}$. Compressive sampling approach can be readily borrowed to find the sparest solution, by formulating the sparse recovery problem as

$$\min_{\boldsymbol{c}} \|\boldsymbol{c}\|_0 \quad \text{subject to} \quad \boldsymbol{\Psi}\boldsymbol{c} = \boldsymbol{u}, \tag{13}$$

where $\|\cdot\|_0$ indicates the ℓ_0-norm, i.e. the number of non-zero terms. However, since ℓ_0-norm is non-convex and discontinuous, the above problem is NP-hard. Therefore, ℓ_0 minimization is usually replaced by its convex relaxation where ℓ_1-norm of the solution $\boldsymbol{c}$ is minimized instead, i.e.

$$\min_{\boldsymbol{c}} \|\boldsymbol{c}\|_1 \quad \text{subject to} \quad \boldsymbol{\Psi}\boldsymbol{c} = \boldsymbol{u}, \tag{14}$$

ℓ_1 minimization is the closest convex problem to ℓ_0 minimization. It has been shown that when $\boldsymbol{\Psi}$ is sufficiently incoherent, as will be explained later, and $\boldsymbol{c}$ is sufficiently sparse, the solution of ℓ_0 minimization is unique and is equal to the solution of ℓ_1 minimization [35]. The minimization in (14) is named basis pursuit [36], and it can be solved using linear programming. If the measurements are known to be noisy, the problem can be reformulated as

$$\min_{\boldsymbol{c}} \|\boldsymbol{c}\|_1 \quad \text{subject to} \quad \|\boldsymbol{\Psi}\boldsymbol{c} - \boldsymbol{u}\|_2 < \epsilon, \tag{15}$$

which is known as the basis pursuit denoising problem. ϵ controls the inaccuracy tolerance and can be prescribed if the distribution of measurement noise is known. For example, if the measurement noise follows a normal distribution with zero mean and standard deviation σ_n, then ϵ is naturally set to be $\sqrt{M\sigma_n^2}$. When instead of measurement samples, simulation samples are used in the PC estimation, and a distribution for numerical or modeling error is absent, one may use cross-validation to determine the tolerance parameter ϵ. This section continues with a brief review on theorems developed on recoverability of compressive sampling and efforts heretofore made to improve the accuracy of PC estimation using compressive sampling.

2.3. PC recoverability using compressive sampling

Whether desired solution can be recovered from compressive sampling or not mainly depends on the properties of the measurement matrix $\boldsymbol{\Psi}$. One of the properties that are shown to be significantly relevant is the Restricted Isometry Property (RIP) or the s-restricted isometry constant for the measurement matrix [37]. Formally speaking the s-restricted isometry constant for a matrix $\boldsymbol{\Psi} \in \mathbb{R}^{M\times K}$ is defined to be the smallest $\delta_s \in (0,1)$ such that

$$(1-\delta_s)\|\boldsymbol{c}\|_2^2 \leqslant \left\|\boldsymbol{\Psi}^{s^*}\boldsymbol{c}\right\|_2^2 \leqslant (1+\delta_s)\|\boldsymbol{c}\|_2^2, \tag{16}$$

for every submatrix $\boldsymbol{\Psi}^{s^*} \in \mathbb{R}^{M\times s^*}, s^* \leqslant s$, of $\boldsymbol{\Psi}$ and every vector $\boldsymbol{c} \in \mathbb{R}^{s^*}$. Thus defined, a small RIP constant for a measurement matrix will lead to a situation wherein for any given sparse signal, energy of the measured signal is not very different from the energy of the given signal itself; hence the appeal of a small RIP constant. The following theorem provides an upperbound on RIP constant so that the optimization problem (15) leads to accurate recovery.

Theorem 1 ([37])**.** *Let $\boldsymbol{\Psi} \in \mathbb{R}^{M\times K}$ with RIP constant δ_{2s} such that $\delta_{2s} < \sqrt{2}-1$. For a given $\bar{\boldsymbol{c}}$, and noisy measurement $\boldsymbol{y} = \boldsymbol{\Psi}\bar{\boldsymbol{c}} + \boldsymbol{\epsilon}$, $\|\boldsymbol{\epsilon}\|_2 \leqslant \epsilon$, let $\boldsymbol{c}$ be the solution of*

$$min \|\boldsymbol{c}\|_1 \quad subject\ to \quad \|\boldsymbol{\Psi}\boldsymbol{c} - \boldsymbol{y}\|_2 < \epsilon. \tag{17}$$

Then the reconstruction error satisfies

$$\|\boldsymbol{c} - \bar{\boldsymbol{c}}\| \leqslant C_1 \frac{\|\boldsymbol{c} - \boldsymbol{c}^*\|}{\sqrt{s}} + C_2\epsilon, \tag{18}$$

where C_1 and C_2 only depend on δ_{2s} and $\boldsymbol{c}^*$ is the vector $\boldsymbol{c}$ with all but the s-largest entries set to be zero. If $\boldsymbol{c}$ is s-sparse and the measurements are noiseless, then the recovery is exact.

Calculating RIP constant for a given matrix is an NP-complete problem [38]. The following theorem gives a probabilistic upperbound on RIP constant for bounded orthonormal systems. Let $\{\psi_n\}_{1\leqslant n\leqslant K}$ be an orthonormal bounded system of functions on $\mathcal{D}$, where $\mathcal{D}$ is endowed with a probability measure ν. Specifically, we have

$$\int_{\mathcal{D}} \psi_n(\xi)\psi_m(\xi)\nu(\xi)\,\mathrm{d}\xi = \delta_{m,n}, \quad 1 \leqslant m, n \leqslant K, \tag{19}$$

and the following uniform bound,

$$\|\psi_n\|_\infty = \sup_{\xi\in\mathcal{D}} |\psi_n(\xi)| \leqslant L, \quad \text{for all} \quad 1 \leqslant n \leqslant K. \tag{20}$$

Consequently, an orthonormal polynomial system $\{\psi_{\boldsymbol{\alpha}}\}_{\boldsymbol{\alpha}\in\Lambda_{d,k}}$ defined in (5) is a bounded orthonormal system

if it is uniformly bounded by

$$\sup_{\boldsymbol{\alpha}\in\Lambda_{d,k}} \|\psi_{\boldsymbol{\alpha}}\|_\infty = \sup_{\boldsymbol{\alpha}\in\Lambda_{d,k}} \sup_{\boldsymbol{\xi}\in\Omega} |\psi_{\boldsymbol{\alpha}}(\boldsymbol{\xi})| \leqslant L, \tag{21}$$

for some constant $L \geq 1$. Hereinafter, let us refer to the bound L as the local-coherence.

Theorem 2 ([39]). *Let $\{\psi_n\}_{1\leqslant n\leqslant K}$ be a bounded orthonormal system satisfying Equations (19) and (20). Also let $\boldsymbol{\Psi} \in \mathbb{R}^{M\times K}$ be a measurement matrix with entries $\left\{\psi_{ij} = \psi_j(\xi^{(i)})\right\}_{1\leqslant i\leqslant M, 1\leqslant j\leqslant K}$, where $\xi^{(1)}, ..., \xi^{(M)}$ are random samples drawn from measure ν. Assuming that*

$$M \geqslant C\delta^{-2}L^2 s \log^3(s)\log(K), \tag{22}$$

then with probability at least $1 - K^{\beta\log^3(s)}$ the RIP constant δ_s of $\frac{1}{\sqrt{M}}\boldsymbol{\Psi}$ satisfies $\delta_s \leqslant \delta$. The $C, \beta > 0$ are universal constants.

It is worthwhile to highlight the differences between these two theorems. Theorem 1 is a deterministic theorem, while Theorem 2 is probabilistic. Theorem 1 provides a deterministic guarantee that is universal, in the sense that once a specific RIP property is satisfied for samples of a class of measurement matrices, then the corresponding error bounds hold for all kinds of target expansions and recovery is always exact for cases with noiseless measurement and s-sparse target. On the other hand, when the condition of Theorem 2 is satisfied, i.e. sufficient number of samples are provided, the recovery accuracy can be achieved with high probability, and not with probability 1.

3. Near-optimal sampling strategy

Suppose that a limited sampling budget allows for M samples to be drawn. Our objective, then, is to optimally select M locations in the parameter space, at which samples of QoI are computationally or experimentally evaluated. Let us pose a slightly different, but closely related, question: out of a sufficiently large pool of M_p candidate sample locations, how can one identify the M "optimal" locations? The sufficiently large number M_p can be determined such that an accurate regression for the given stochastic system is achieved. In what follows, we first discuss the sampling strategies that aim at improving the local-coherence.This will then be followed by discussions about other relevant properties, in particular those of the measurement matrix, that have implications on accuracy and stability of sparse recovery, and how a 'near-optimal' sampling strategy can be designed to improve these relevant properties.

3.1. Sampling strategies focusing on local-coherence

As is evident by the theoretical results of Section 2.3, specifically in Theorem 2, local-coherence has an implication on recovery accuracy. This has motivated researchers to develop new sampling strategies such that this property is improved. The conventional sampling approach, namely the standard (random) sampling approach, involves drawing random realizations from uniform and normal distributions for Legendre and Hermite PC expansions, respectively. In [24], it was recommended to use a preconditioned Legendre measurement matrix, where samples are taken from Chebyshev distribution, instead of a uniform distribution, to improve accuracy. With respect to this new sampling distribution, new orthogonal bases were calculated by scaling or weighting the Legendre polynomials. It was shown in [33] that when system dimensionality d is larger than polynomial order k, the Chebyshev preconditioning approach will result in a polynomial system with a larger local-coherence, L, and subsequently in less accurate results compared with that from standard sampling. An alternative sampling strategy was proposed in [25], where the Legendre polynomial system was first turned into a discrete orthogonal system using Gauss quadrature points, from which samples are drawn randomly for sparse recovery. But, it was shown that the new discretized orthogonal system improves the bound L only when $k \geqslant d$, and for $k < d$, recovery accuracy will be lowered than that from standard sampling.

In [26], a coherence-optimal, or to be compliant with the terminology in this paper, a local-coherence-optimal sampling approach was introduced, where instead of directly sampling from the probability measure, $\rho(\boldsymbol{\xi})$, samples are drawn from a different "optimal" probability measure, $\rho_o(\boldsymbol{\xi})$, constructed according to

$$\rho_o(\boldsymbol{\xi}) = C^2\rho(\boldsymbol{\xi})B^2(\boldsymbol{\xi}), \tag{23}$$

where C is a normalizing constant and

$$B(\boldsymbol{\xi}) := \max_{\boldsymbol{\alpha}\in\Lambda_{d,k}} |\psi_{\boldsymbol{\alpha}}(\boldsymbol{\xi})|\,. \tag{24}$$

Corresponding to this probability measure, the weight function, used to retrieve the orthogonality of the bases, should be

$$w(\boldsymbol{\xi}) = \frac{1}{B(\boldsymbol{\xi})}. \tag{25}$$

Accordingly, the ℓ_1-minimizations in (14) and (15) were adjusted to include the weight functions. The following weighted ℓ_1-minimizations were then used for sparse recovery

$$\min_{\boldsymbol{c}} \|\boldsymbol{c}\|_1 \quad \text{subject to} \quad \boldsymbol{W}\boldsymbol{\Psi}\boldsymbol{c} = \boldsymbol{W}\boldsymbol{u}, \tag{26}$$

$$\min_{\boldsymbol{c}} \|\boldsymbol{c}\|_1 \quad \text{subject to} \quad \|\boldsymbol{W}\boldsymbol{\Psi}\boldsymbol{c} - \boldsymbol{W}\boldsymbol{u}\|_2 < \epsilon, \tag{27}$$

where $\boldsymbol{W}$ is the $M \times M$ diagonal weight matrix, with $\boldsymbol{W}(i,i) = w(\boldsymbol{\xi}^{(i)})$ for $i = 1, \cdots, M$. It was also proved in [26] that this sampling approach results in the lowest local-coherence among all sampling approaches, and can thus outperform standard sampling for both cases, i.e., $k < d$ and $k \geqslant d$.

In what follows, we discuss how this strategy can be further improved by considering other relevant properties with implications on recovery accuracy.

3.2. Sampling strategies focusing on measurement matrix properties

It is obvious that to recover all s-sparse vectors $\bar{\boldsymbol{c}}$ from observation vector $\boldsymbol{u} = \boldsymbol{\Psi}\bar{\boldsymbol{c}}$, using the ℓ_0 minimization problem in (13), we must have $\boldsymbol{\Psi}\bar{\boldsymbol{c}} \neq \boldsymbol{\Psi}\boldsymbol{c}'$ for any pair of distinct s-sparse vectors $\boldsymbol{c}', \bar{\boldsymbol{c}}$. In other words, we want $\boldsymbol{c}' \neq \bar{\boldsymbol{c}}$ to lead to distinct observation vectors. For this to happen, the measurement matrix should not satisfy $\boldsymbol{\Psi}(\bar{\boldsymbol{c}} - \boldsymbol{c}') = 0$, for any vector $\bar{\boldsymbol{c}} - \boldsymbol{c}'$ that has $2s$ or fewer nonzero terms. This means that the null space of $\boldsymbol{\Psi}$ should contain no vector with $2s$ or fewer nonzero terms. In compressive sampling, this property is mostly characterized using *spark* [41]. The spark of a matrix is the smallest number of its columns that are linearly dependent. Using the definition of spark, it is then straightforward to guarantee that the solution of ℓ_0 minimization given in (13) exactly recovers the s-sparse $\bar{\boldsymbol{c}}$ if $\text{spark}(\boldsymbol{\Psi}) > 2s$ [42]. Moreover, it can be shown that in this case the ℓ_1 minimization in (14) also recovers the exact solution, $\bar{\boldsymbol{c}}$ [43].

Clearly, a measurement matrix with a larger spark allows exact recovery using compressive sampling for a larger set of target signals (solution vectors). Therefore, it is desired to maximize the spark of measurement matrices. However, computing the spark of a matrix is an NP-hard problem [44]. As an alternative, one can analyze recovery guarantees using alternative properties which are easier to compute. One such property is the mutual-coherence, which for a given matrix $\boldsymbol{\Psi} \in \mathbb{R}^{M\times K}$ is defined as the maximum absolute normalized inner product, i.e. cross-correlation, between its columns [18, 45]. Let $\psi_1, \psi_2, ..., \psi_K \in \mathbb{R}^M$ be the columns of matrix $\boldsymbol{\Psi}$. The mutual-coherence of matrix $\boldsymbol{\Psi}$, denoted by $\mu(\boldsymbol{\Psi})$, is then given by

$$\mu(\boldsymbol{\Psi}) := \max_{1\leqslant i,j\leqslant K, i\neq j} \frac{|\psi_j^T \psi_i|}{\|\psi_j\|_2 \|\psi_i\|_2}. \tag{28}$$

The following simple proposition explains how spark and mutual coherence are related.

Proposition 1 ([46])**.** For any matrix $\boldsymbol{\Psi} \in \mathbb{R}^{M\times K}$ the following holds:

$$\text{spark}(\boldsymbol{\Psi}) \geqslant 1 + \frac{1}{\mu(\boldsymbol{\Psi})}. \tag{29}$$

Mutual-coherence is an indicator of the worst interdependence, i.e. maximum cross-correlation between the columns of a matrix. It is zero for an orthogonal matrix and is strictly positive when $M < K$. Proposition 1 makes it obvious that a small value for mutual-coherence is desired. It may be concluded that measurement matrix $\boldsymbol{\Psi}$ should be designed in a way that its mutual-coherence, i.e. its maximum cross-correlation, is minimized. However, it has been observed that minimizing maximum cross-correlation does not necessarily improve the recovery accuracy of compressive sampling [43, 47]. This is because Equation (29) only provides

a lower bound on the spark. Therefore, minimizing mutual-coherence considers only the worst-case scenario and fails to account for other possibilities for improving compressive sampling performance [47].

Our objective in this work is to design a measurement matrix $\mathbf{\Psi}$ which leads to better compressive sampling performance *on average*, and not merely in the worst-case scenario. To this end, we need to determine which property of the measurement matrix should be optimized. This property should ideally be one that directly and sufficiently control the accuracy of compressive sampling. A widely-accepted property is the RIP constant, as suggested by Theorem 1. As previously mentioned, calculating RIP constant for a given matrix is an NP-complete problem. However, it is known that the eigenvalues of each column-submatrix $\mathbf{\Psi}^{s^*} \in \mathbb{R}^{M \times s^*}, s^* \leqslant s$ can be bounded by the RIP constant, δ_s

$$1 - \delta_s \leqslant \lambda_{\min}(\mathbf{\Psi}^{s^*T}\mathbf{\Psi}^{s^*}) \leqslant \lambda_{\max}(\mathbf{\Psi}^{s^*T}\mathbf{\Psi}^{s^*}) \leqslant 1 + \delta_s. \tag{30}$$

Therefore, one may suggest to minimize the condition number of all column-submatrices with s or fewer columns. The challenge, however, is that the sparsity level of solution, s, is not known in advance. Moreover, calculating such a combinatorial measure is not a trivial task and can be computationally impossible [47].

As a result, establishing a single matrix property that sufficiently guarantees the accuracy of compressive sampling method and is easily computable still remains an open challenge [43, 47]. To sidestep this challenge, efforts have focused on identifying properties or measures than are relatively better than maximum cross-correlation, or mutual-coherence. In [47], for the first time, it was suggested that a t-averaged mutual-coherence be minimized. Denoted by $\mu_t(\mathbf{\Psi})$, the t-averaged mutual-coherence for a measurement matrix $\mathbf{\Psi}$ is defined as the average of cross-correlations larger than threshold t, i.e.

$$\mu_t(\mathbf{\Psi}) = \frac{\sum\limits_{1 \leqslant i,j \leqslant K, i \neq j} \mathbb{1}(|g_{ij}| \geqslant t).\,|g_{ij}|}{\sum\limits_{1 \leqslant i,j \leqslant K, i \neq j} \mathbb{1}(|g_{ij}| \geqslant t)}, \tag{31}$$

where $\mathbb{1}(\cdot)$ is the indicator function, g_{ij} is the ijth component of the Gram matrix $\boldsymbol{G} = \tilde{\mathbf{\Psi}}^T\tilde{\mathbf{\Psi}}$, and $\tilde{\mathbf{\Psi}}$ is the column-normalized version of $\mathbf{\Psi}$. It has been shown that recovery accuracy can be significantly improved if instead of a random measurement matrix, the measurement matrix is optimized based on $\mu_t(\mathbf{\Psi})$. However, it was later shown in [48] that a $\mu_t(\mathbf{\Psi})$-optimized measurement matrix is not robust when target signals are attached to some noise, and as such, not exactly sparse.

In order to improve this robustness, efforts have focused on optimizing measurement matrices by considering all the cross-correlations, and not the ones larger than a threshold. Specifically, these efforts seek to minimize the distance between the Gram matrix, $\boldsymbol{G}$, and the corresponding identity matrix. With $\|\cdot\|_F$ denoting the Frobenius norm, it has been shown that the measurement matrix determined by solving the following minimization problem

$$\min_{\mathbf{\Psi} \in R^{M \times K}} \left\| I_K - \tilde{\mathbf{\Psi}}^T\tilde{\mathbf{\Psi}} \right\|_F^2, \tag{32}$$

can lead to a signal recovery that is both more accurate and more robust compared with that using the $\mu_t(\mathbf{\Psi})$-optimized matrix [43, 49, 50, 51]. In what follows, we propose our near-optimal sampling strategy which considers measures of maximum and average cross-correlation of measurement matrices, and optimizes sample locations accordingly.

3.3. Near-optimal sampling strategy

To achieve a better robustness, as argued earlier, we aim to select sample locations that collectively minimize a measure of average cross-correlation given by

$$\gamma(\mathbf{\Psi}) := \frac{1}{N} \left\| I_K - \tilde{\mathbf{\Psi}}^T\tilde{\mathbf{\Psi}} \right\|_F^2, \tag{33}$$

where $N := K \times (K-1)$ is the total number of column pairs. For the sake of brevity, we refer to $\gamma(\mathbf{\Psi})$ as the *average cross-correlation*, even though precisely speaking, it is the average of squares of cross-correlations. It should be noted that minimizing this average cross-correlation alone will not necessarily result in the smallest maximum cross-correlation, and the mutual-coherence could still be undesirably large and the sparse recovery significantly inaccurate. As a remedy, we seek to minimize both the average cross-correlation $\gamma(\mathbf{\Psi})$ and the mutual-coherence $\mu(\mathbf{\Psi})$, simultaneously in a multi-objective problem.

To solve the optimization problem in (32), an option is to adopt the solution approaches proposed for (32). These are mainly iterative approaches based on gradient descent method [49] or bound-optimization method [50]. However, these algorithms are applicable for measurement matrices that are less restricted compare to PC measurement matrix, in that their components are not constrained to be evaluations of specific orthogonal bases at certain sample locations, and as such can be freely optimized. Therefore, a greedy algorithm is the natural option for solving . Our greedy algorithm for near-optimal sampling strategy begins by populating a large pool of candidate sample locations in the multidimensional parameter space, and seeks to find the near-optimal M locations and the corresponding $M \times K$ measurement matrix. We select the first sample location randomly from the candidates pool, and this constitutes the first row in the measurement matrix. At each step of the algorithm, a new sample location, or matrix row is added. This is done by searching through the candidates pool for the best sample location. The best new location is one that produces a matrix with the smallest average cross-correlation and smallest mutual-coherence.

Let $\mathbf{\Psi}^{\text{pool}}$ denote the $M_p \times K$ measurement matrix associated with the large pool of M_p candidate locations, and $\mathbf{\Psi}^{\text{opt}[M]}$ the near-optimal $M \times K$ row-submatrix of the "pool" matrix $\mathbf{\Psi}^{\text{pool}}$. This submatrix is identified in the algorithm through incremental row-concatenation, as shown in Algorithm 1. In this pseudo-code, $\mathbf{\Psi}^{\text{opt}[i]}$ of size $i \times K$ denotes the near-optimal submatrix at the ith step of the algorithm, and $\mathbf{\Psi}^{\text{pool}}_{(j)}$ represents the jth row in the "pool" matrix.

Algorithm 1 Greedy algorithm for near-optimal sample locations

Initiate $\mathbf{\Psi}^{\text{opt}[1]}$ to be a random row in $\mathbf{\Psi}^{\text{pool}}$
for $i = 2 : M$ **do**
  $\mathbf{\Psi}^{\text{opt}[i]} = \text{row-concatenate}(\mathbf{\Psi}^{\text{opt}[i-1]}, \mathbf{\Psi}^{\text{pool}}_{(1)})$
  for $j = 2 : M_p$ **do**
    $\mathbf{\Psi}^{\text{temp}} = \text{row-concatenate}(\mathbf{\Psi}^{\text{opt}[i-1]}, \mathbf{\Psi}^{\text{pool}}_{(j)})$
    if $\mu(\mathbf{\Psi}^{\text{temp}}) < \mu(\mathbf{\Psi}^{\text{opt}[i]})$ and $\gamma(\mathbf{\Psi}^{\text{temp}}) < \gamma(\mathbf{\Psi}^{\text{opt}[i]})$ **then**
      $\mathbf{\Psi}^{\text{opt}[i]} = \mathbf{\Psi}^{\text{temp}}$
    end if
  end for
end for

As discussed in Section 2.3, it has been shown that local-coherence-optimal sampling improves the PC coefficients recovery accuracy, over standard sampling. To exploit this advantage, we use the local-coherence-optimal sampling strategy to generate the large pool of candidate sample locations. In order to retain orthogonality, we just need to substitute $\mathbf{\Psi}^{\text{pool}}$ with $\boldsymbol{W}\mathbf{\Psi}^{\text{pool}}$ in Algorithm 1, where $\boldsymbol{W}$ is the diagonal weight matrix defined earlier in Section 3.1. Compared to standard sampling and local-coherence-optimal sampling, our sampling strategy incurs extra computational cost due to additional row selections from the candidates pool in the greedy algorithm. However, this additional cost is typically negligible with respect to the cost of sample collection, especially when the benefit of improved accuracy is also considered.

We refer to the proposed sampling approach as near-optimal sampling strategy. The reason we do not use the term 'optimal' is two-fold: First, even though studies in signal processing have shown significant improvement in recovery accuracy by minimizing (33), this orthogonality property does not by itself establish a sufficient criterion for recovery accuracy of compressive sampling (A measure that is both sufficient and tractable has not been identified in the literature.) Second, our approach can be sensitive to (i) the random choice of candidate locations and (ii) the random choice of the initial location (first row in the submatrix), and is therefore not fully deterministic, and as such not optimal. In the next section, we provide numerical examples to demonstrate advantages of our proposed sampling strategy.

4. Numerical examples

To demonstrate the advantage of the near-optimal sampling strategy over other sampling approaches, four target functions are considered in this section: (i) a low-dimensional high-order polynomial function, (ii) a high-dimensional low-order polynomial function, (iii) a six-dimensional generalized Rosenbrock function, and (iv) the solution to a stochastic diffusion problem. This would allow for a comprehensive comparison of the sampling strategies for a variety of models with different combinations of dimension and order. The first three target functions are chosen to be exactly s-sparse (or sparse, in short), whereas the last one is

approximately s-sparse (or compressible, in short). In all examples, optimization problems (14) and (26), corresponding to a noise-less measurement setting, were formed and solved using the SPGL1 package [52]. For the sake of brevity, we only report the results for Legendre polynomial expansions in this paper, but note that similar improvements in recovery accuracy were observed when near-optimal sampling was applied to Hermite polynomial expansions. In all these examples, the proposed near-optimal sampling approach is compared with (i) standard sampling, and (ii) local-coherence-optimal sampling, or in short, 'coherence-optimal' sampling.

4.1. Low-dimensional high-order sparse PCE

Let us consider the target function, $u(\boldsymbol{\Xi})$, to be a sparse 20th-order Legendre polynomial expansion in a two-dimension random space with uniform density on $[-1,1]^2$, manufactured according to

$$u(\boldsymbol{\Xi}) = \sum_{i=1}^{5} \Xi_1^{2i} \Xi_2^{2i}. \tag{34}$$

Using limited samples from this exact function, the objective is to recover the sparse coefficient vector. In selecting the sample locations to form the measurement matrix and data vector, the proposed near-optimal sampling approach is compared versus random sampling strategies, namely (i) standard sampling and (ii) coherence-optimal sampling. The candidates pool in the near-optimal sampling approach includes $100,000$ coherence-optimal samples. For all three approaches, we report the performance results obtained by 100 independent runs. This is to capture the variability induced by small sample size in standard and coherence-optimal sampling approaches. In the near-optimal approach, these independent runs allow us to account for (i) the sensitivity of performance with respect to the choice of initial sample location, and (ii) the variability induced by the finite size of the candidate pool, and (iii) the variability induced by the fact that our cross-correlation measures do not necessarily guarantee CS recovery accuracy.

Figure 1a shows the median, and 1st and 3rd quantiles of relative ℓ_2 error and Figure 1b shows the mean of relative ℓ_2 error. The relative ℓ_2 error is calculated as $\|\boldsymbol{c} - \bar{\boldsymbol{c}}\|_2 / \|\bar{\boldsymbol{c}}\|_2$, where $\bar{\boldsymbol{c}}$ is the exact coefficient vector and $\boldsymbol{c}$ is the solution of ℓ_1 minimization in (14). As is apparent in Figure 1a and 1b , coherence optimal sampling results in a smaller error, compared to standard sampling, as the samples are drawn from a distribution with a smaller bound L defined in (21). However, our proposed sampling outperforms the coherence-optimal one. To explain this, Figures 1c and 1d show the median, and 1st and 3rd quantiles of mutual-coherence and average cross-correlation for the measurement matrix, respectively. It can be seen that near-optimal sampling beats the other two approaches in both measures, leading to the higher observed accuracy in the PCE estimation.

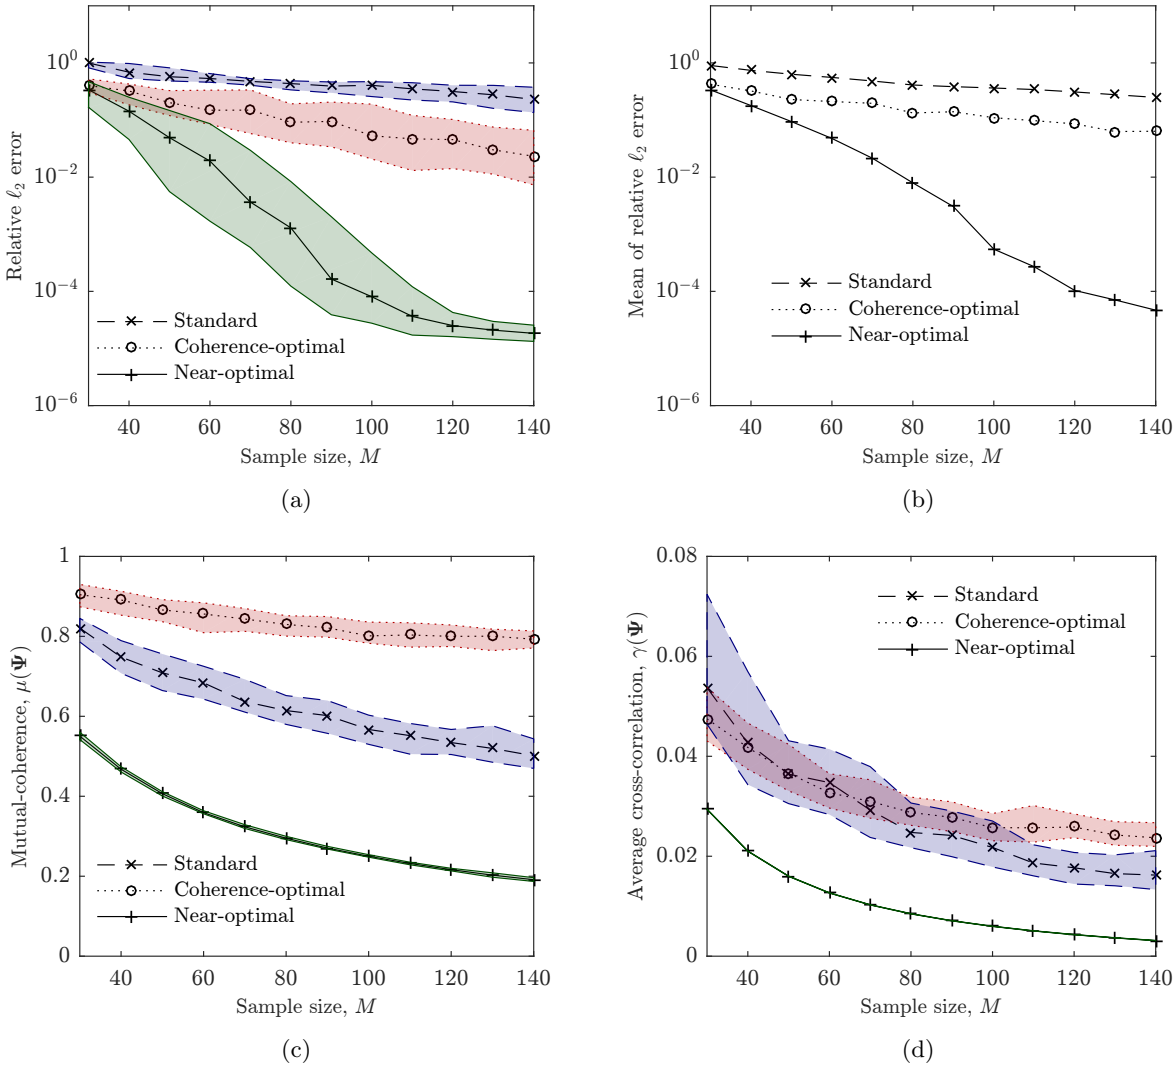

Figure 1: Comparison of near-optimal sampling results with standard and coherence-optimal sampling methods for the low-dimensional high-degree manufactured PCE of Equation (34). The following measures are compared: (a) median, and 1st and 3rd quantiles of relative ℓ_2 error (b) mean of relative ℓ_2 error, (c) median, and 1st and 3rd quantiles of mutual-coherence, and (d) median, 1st and 3rd quantiles of average cross-correlation.

4.2. High-dimensional low-order sparse PCE

As a contrasting example, let us consider the target function, $u(\boldsymbol{\Xi})$, to be a sparse second order Legendre polynomial expansion in a 20-dimensional random space with uniform density on $[-1,1]^{20}$, manufactured according to

$$u(\boldsymbol{\Xi}) = \sum_{i=1}^{19} \Xi_i \Xi_{i+1}. \tag{35}$$

To compare the performance of near-optimal sampling with standard and coherence optimal sampling on this target function, the numerical results were obtained under a setting similar to that in the previous example: 100 independent runs were used for all three approaches, and 100,000 coherence-optimal samples constituted the initial sample pool in the near-optimal approach of Algorithm 1. Figure 2a shows the median, and 1st and 3rd quantiles of relative ℓ_2 error. Figure 2b shows the mean of relative ℓ_2 error. Figure 2c and 2d show the improvement in the median, and 1st and 3rd quantiles of mutual-coherence and average cross-correlation, respectively, when the near-optimal sampling is used, which has translated into the observed improvement in recovery accuracy.

It should be highlighted that standard sampling is mostly suitable for high-dimensional and low-order cases such as this example [18]. It should also be noted that in the examples in this section and Section

4.1, the measurement matrices have the same number of columns ($K = 231$); however, as it can be seen in Figures 2c and 2d, the high-dimensionality of this example has resulted in smaller mutual-coherences and average cross-correlation for various sample sizes. Consequently, we expect to observe that recently developed sampling strategies do not outperform standard sampling in these high-dimensional low-order problem cases [25, 53]. However, as it can be seen in figures 2a and 2b, the near-optimal sampling approach does outperform both standard and coherence optimal sampling. Although, this improvement in terms of mean or median of relative error may not seem significant, it should be noted that this is a rather "extreme" case, where not only the response is high-dimensional, but its order is also very low. This could be thought of as a lower bound for the improvement offered by the near-optimal sampling approach. That is, the fact that the near-optimal sampling strategy still shows improvement, even non-significant, for such extreme cases where standard sampling is supposed to work well, can be considered as a numerical evidence that it will most likely outperform other sampling strategies in all the other cases.

As we already discussed in Section 3.2, in order to have exact recovery we need the spark of measurement matrix to be larger than two times the number of non-zero coefficients. At small sample sizes none of the sampling approaches meet this requirement. Therefore, as it can be seen in figures 2a and 2b, at low sample sizes, all three sampling strategies recover poor approximations. As the number of sample increases the chance to meet this requirement also increases; hence the variability in recovery accuracy also increases. In near-optimal sampling approach we select sample locations such that the orthogonality of measurement matrix is improved, thereby enhancing the probability of achieving a larger spark and consequently exact recovery. To further clarify this, we demonstrate the advantage of near-optimal sampling by comparing another performance measure, namely the success rate. The success rate, here, is defined as the ratio of trials, out of the total 100 trials, that result in a relative error smaller than 10^{-7}. Figure 2e shows that near-optimal sampling consistently results in a higher success rate, and is thus expected to produce accurate recoveries, more frequently.

4.3. Generalized Rosenbrock function

Compared to the previous two examples with "extreme" dimension and order combinations, the third target function can be chosen to be of "moderate" combination of dimension and order. Specifically, let us consider the following 6-dimensional generalized Rosenbrock function with random inputs following a uniform density on $[-1, 1]^6$,

$$u(\boldsymbol{\Xi}) = \sum_{i=1}^{5} 100(\Xi_{i+1} - \Xi_i^2)^2 + (1 - \Xi_i)^2. \tag{36}$$

The objective is to recover the corresponding sparse Legendre polynomial expansion. Similar setting is used for comparison of the three sampling strategies: 100 independent runs for all sampling approaches, together with a candidate pool of 100,000 coherence-optimal samples for near-optimal sampling of Algorithm 1. Figure 3a demonstrates the improvements in the median, and 1st and 3rd quantiles of relative ℓ_2 error when near-optimal sampling is used. Figure 3b demonstrates the improvements in terms of the mean of relative ℓ_2 error. Figure 3c and 3d show that near-optimal sampling results in smaller mutual-coherence and average cross-correlation, in terms of their median, and 1st and 3rd quantiles. Similar to the previous example, we note that at small sample sizes all approaches result in equally inaccurate recoveries as all approaches fail to achieve a measurement matrix with sufficiently large spark, i.e. $\text{spark}(\boldsymbol{\Psi}) > 2s$. A larger sample sizes results in a larger lower bound for the spark of measurement matrix and a higher probability to achieve the critical spark value. Therefore, we observe more variability in recovery accuracy as the sample size increases. Figure 3e shows that using near-optimal sampling leads to significant improvement in the success rate, i.e. the ratio of trials with relative errors smaller than 10^{-7}. This is a direct result of improving the orthogonality of measurement matrix, i.e. enhancing the chance of achieving critical spark valu0.48e in near-optimal sampling.

0.47

4.4. Stochastic diffusion problem

In this example we consider a stochastic diffusion problem in a one dimensional physical domain, given by

$$\begin{aligned} -\frac{\partial}{\partial x}\left(a(x, \boldsymbol{\Xi})\frac{\partial u}{\partial x}(x, \boldsymbol{\Xi})\right) &= 2, \quad x \in (0, 1), \\ u(0, \boldsymbol{\Xi}) = 0, \quad u(1, \boldsymbol{\Xi}) &= 0, \quad \boldsymbol{\Xi} \in [-1, 1]^{10}. \end{aligned} \tag{37}$$

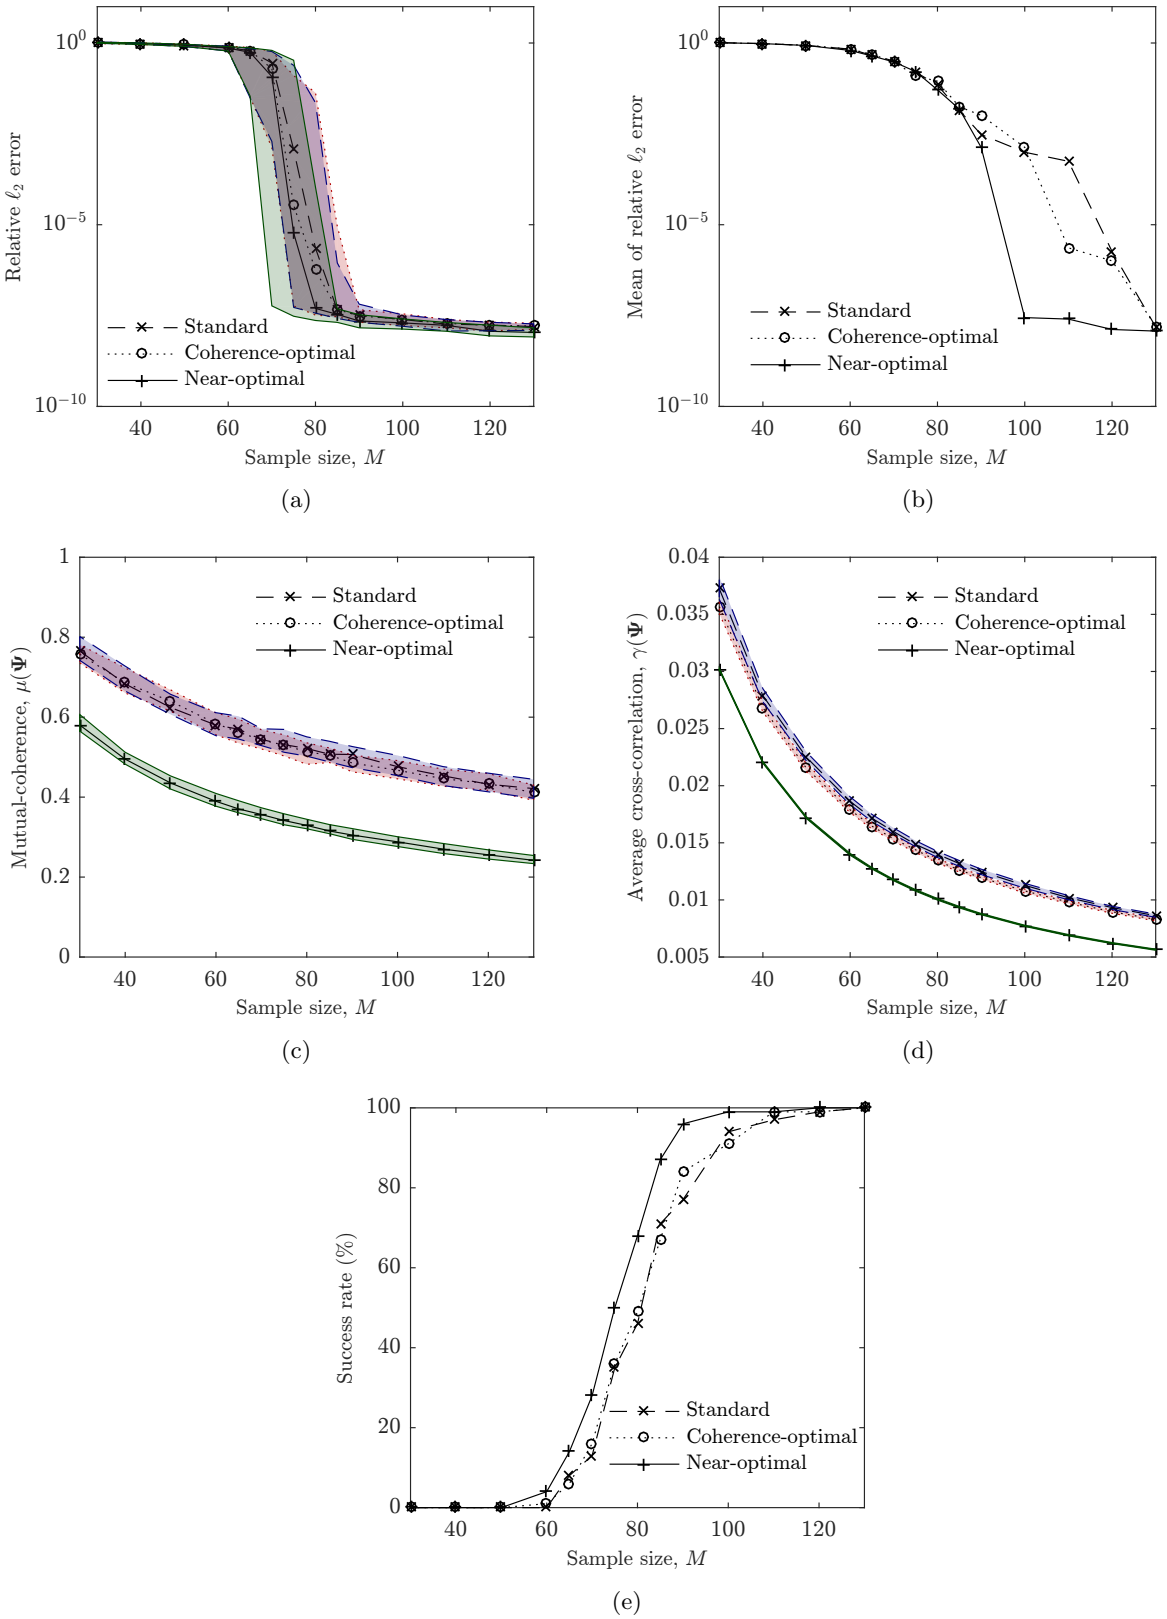

Figure 2: Comparison of near-optimal sampling results with standard and coherence-optimal sampling methods for the high-dimensional low-degree manufactured PCE of Equation (35). The following measures are compared: (a) median, and 1st and 3rd quantiles of relative ℓ_2 error (b) mean of relative ℓ_2 error, (c) median, and 1st and 3rd quantiles of mutual-coherence, (d) median, 1st and 3rd quantiles of average cross-correlation, and (e) success rate, defined as the ratio of trials that result in a relative error smaller than 10^{-7}.

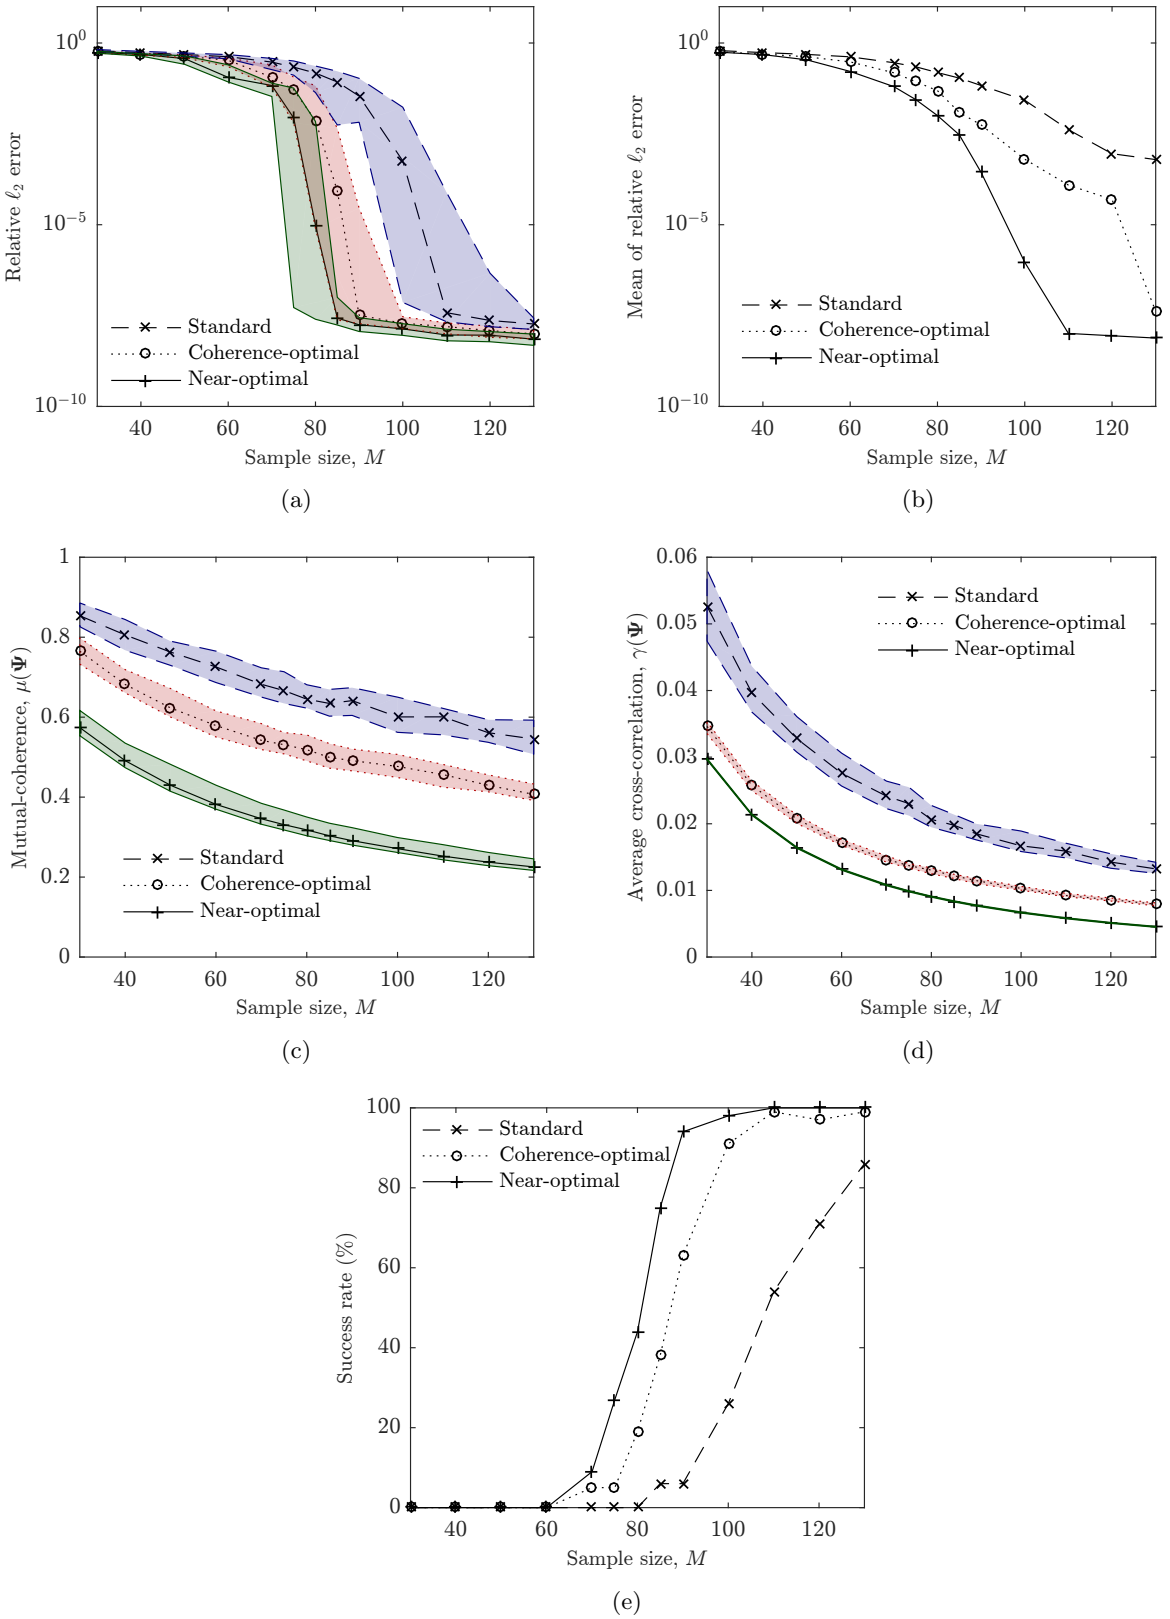

Figure 3: Comparison of near-optimal sampling results with standard and coherence-optimal sampling methods for the generalized Rosenbrock function in Equation (36). The following measures are compared: (a) median, and 1st and 3rd quantiles of relative ℓ_2 error (b) mean of relative ℓ_2 error, (c) median, and 1st and 3rd quantiles of mutual-coherence, (d) median, 1st and 3rd quantiles of average cross-correlation, and (e) success rate, defined as the ratio of trials that result in a relative error smaller than 10^{-7}.

We assume that the diffusion coefficient, $a(x, \boldsymbol{\Xi})$, takes the following analytical form

$$a(x, \boldsymbol{\Xi}) = 1 + \sum_{k=1}^{10} \frac{1}{k^2\pi^2} \cos(2\pi k x)\Xi_k. \tag{38}$$

We consider $\boldsymbol{\Xi}$ to be uniformly distributed on $[-1, 1]^{10}$ and consider the solution of diffusion problem at $u(0.5, \boldsymbol{\Xi})$ to be the quantity of interest. We use 3rd-order Legendre polynomial expansion to approximate the QoI and employ the three sampling approaches to estimate the coefficients of expansion. Similar setting to previous examples is used here: 100 independent runs for all sampling approaches, together with a candidate pool of 100,000 coherence-optimal samples for near-optimal sampling of Algorithm 1. For this example, we define the relative error to be $(\|\boldsymbol{u} - \bar{\boldsymbol{u}}\|_2)/(\|\bar{\boldsymbol{u}}\|_2)$, where $\bar{\boldsymbol{u}}$ is the vector including the exact solutions calculated at 1000 new random samples and $\boldsymbol{u}$ includes the approximate solutions calculated by evaluating the PC expansion at the same 1000 samples. Figure 4a demonstrates the improvement in the median, and 1st and 3rd quantiles of relative ℓ_2 error when near-optimal sampling is used. Figure 4b shows the improvement in the mean of relative ℓ_2 error. The mutual-coherence and average cross-correlation are also compared in figures 4c and 4d, respectively, in terms of their medians, and 1st and 3rd quantiles. To compare the success rates, since the target expansion is not exactly sparse, we consider a looser definition for successful recovery. This is justified by noting that the error levels in Figure 4a are relatively larger than those those in Figures 2a and 3a. Accordingly, we define a recovery to be successful when its relative ℓ_2 error is smaller that 10^{-4}, and show the resulting success rates in Figure 4e. These results show that success rates can be improved significantly by using near-optimal sampling.

5. Conclusion

In this paper, we presented a new sampling strategy, or a design-of-experiment technique, for selecting sample locations for sparse estimation of PCEs using compressive sampling. The sample locations are selected such that (i) the local-coherence property is improved, and (ii) the resulting measurement matrix has the smallest mutual-coherence and the smallest average cross-correlation between its columns. It was discussed how the latter two measures have implications on recovery accuracy and numerical results were presented to support it. A greedy algorithm was introduced that selects a prescribed number of near-optimal locations out of a large pool of candidate locations. The resulting measurement matrix is claimed to be near-optimal, rather than optimal, because the two aforementioned cross-correlation measures may not be the only measures that control the recovery accuracy, and therefore minimizing them does not guarantee optimal performance. Another reason why our algorithm can be sub-optimal is that the greedy algorithm is inherently a non-deterministic algorithm, and as such is susceptible to variabilities induced by random choice of location candidates and initial (first) sample location in the algorithm. Four numerical examples with various combinations of dimensionality and order demonstrated the advantages of our proposed sampling strategy over other sampling strategies, in terms of accuracy and robustness.

6. References

References

[1] J. S. Liu, Monte Carlo strategies in scientific computing, Springer Science & Business Media, 2008.

[2] R. E. Caflisch, Monte Carlo and quasi-Monte Carlo methods, Acta numerica 7 (1998) 1–49.

[3] S. Heinrich, Multilevel Monte Carlo methods, in: International Conference on Large-Scale Scientific Computing, Springer, 2001, pp. 58–67.

[4] R. G. Ghanem, P. D. Spanos, Stochastic finite elements: a spectral approach, Courier Corporation, 2003.

[5] D. Xiu, G. E. Karniadakis, The Wiener–Askey polynomial chaos for stochastic differential equations, SIAM journal on scientific computing 24 (2) (2002) 619–644.

[6] M. K. Deb, I. M. Babuška, J. T. Oden, Solution of stochastic partial differential equations using Galerkin finite element techniques, Computer Methods in Applied Mechanics and Engineering 190 (48) (2001) 6359–6372.

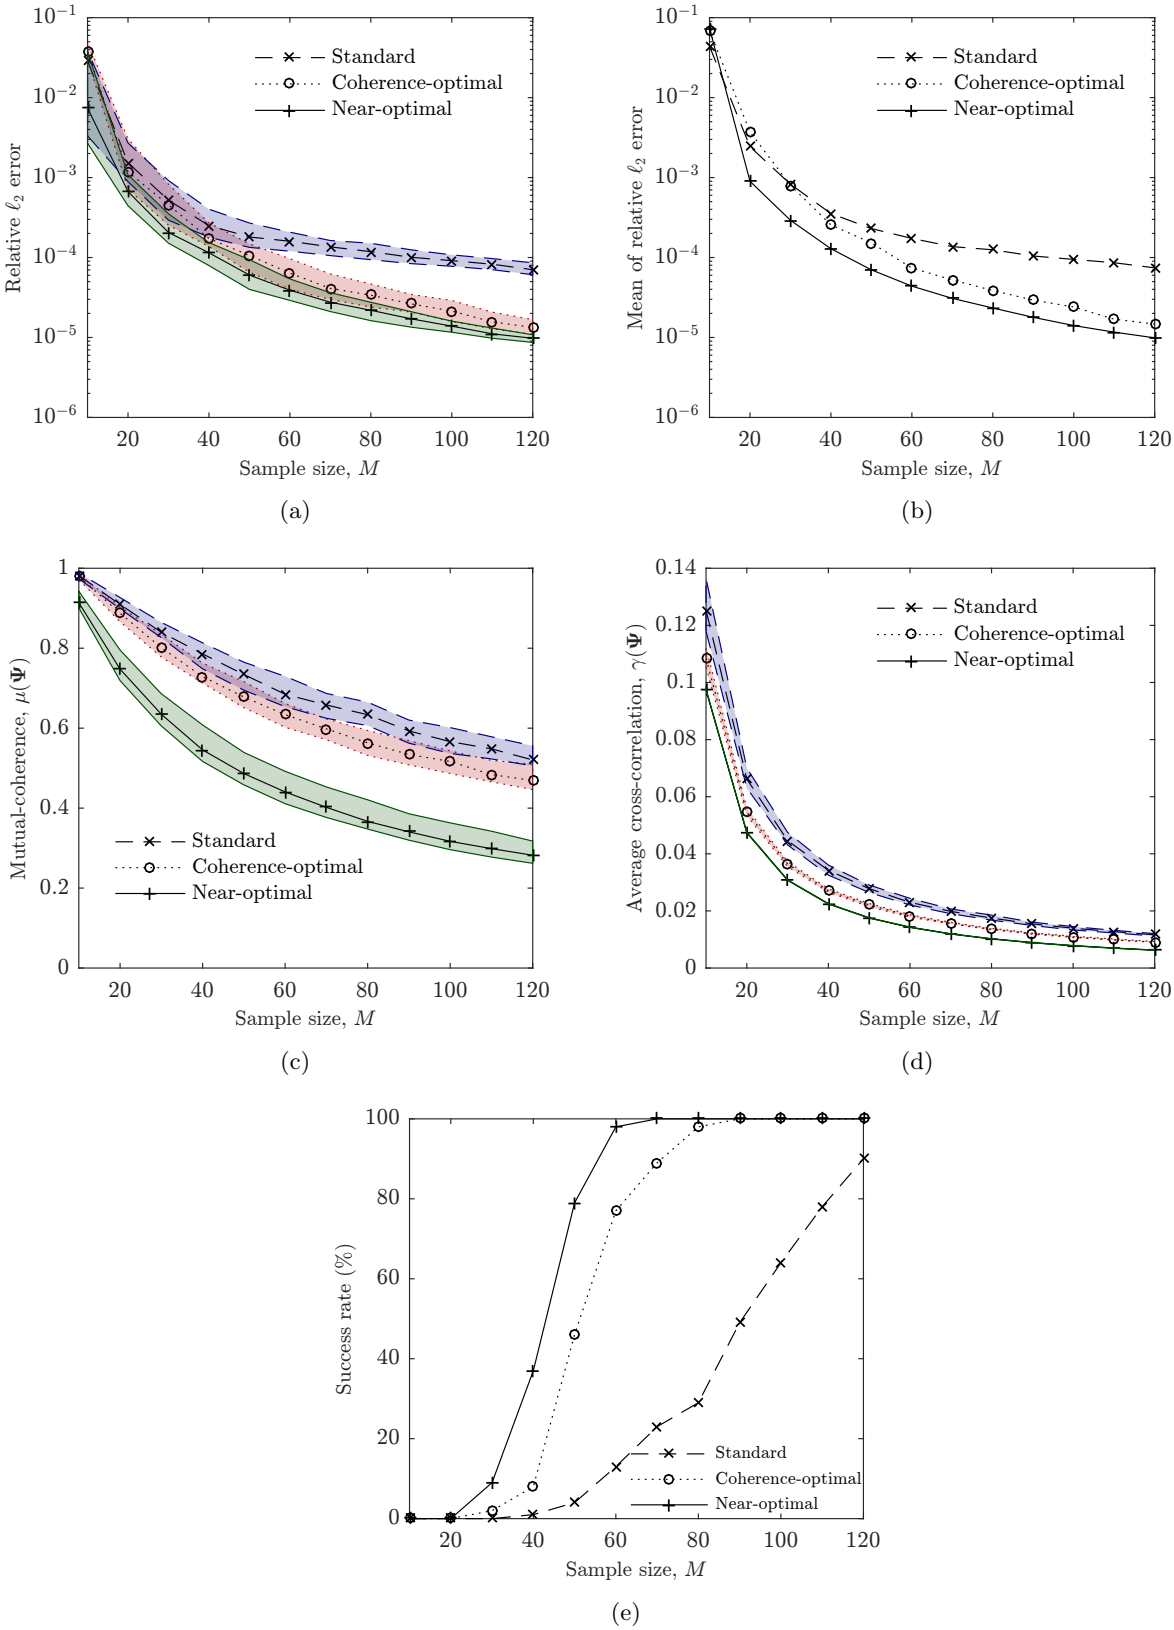

Figure 4: Comparison of near-optimal sampling results with standard and coherence-optimal sampling methods for the solution of the stochastic diffusion problem of Equation (37). The following measures are compared: (a) median, and 1st and 3rd quantiles of relative ℓ_2 error (b) mean of relative ℓ_2 error, (c) median, and 1st and 3rd quantiles of mutual-coherence, (d) median, 1st and 3rd quantiles of average cross-correlation, and (e) success rate, defined as the ratio of trials that result in a relative error smaller than 10^{-4}.

[7] I. Babuska, R. Tempone, G. E. Zouraris, Galerkin finite element approximations of stochastic elliptic partial differential equations, SIAM Journal on Numerical Analysis 42 (2) (2004) 800–825.

[8] M. S. Eldred, C. G. Webster, P. Constantine, Evaluation of non-intrusive approaches for Wiener-Askey generalized polynomial chaos, in: Proceedings of the 10th AIAA Non-Deterministic Approaches Conference, number AIAA-2008-1892, Schaumburg, IL, Vol. 117, 2008, p. 189.

[9] E. Novak, K. Ritter, The curse of dimension and a universal method for numerical integration, in: Multivariate approximation and splines, Springer, 1997, pp. 177–187.

[10] B. Ganapathysubramanian, N. Zabaras, Sparse grid collocation schemes for stochastic natural convection problems, Journal of Computational Physics 225 (1) (2007) 652–685.

[11] D. Xiu, J. S. Hesthaven, High-order collocation methods for differential equations with random inputs, SIAM Journal on Scientific Computing 27 (3) (2005) 1118–1139.

[12] F. Nobile, R. Tempone, C. G. Webster, A sparse grid stochastic collocation method for partial differential equations with random input data, SIAM Journal on Numerical Analysis 46 (5) (2008) 2309–2345.

[13] M. Hegland, Adaptive sparse grids, Anziam Journal 44 (2003) 335–353.

[14] X. Ma, N. Zabaras, An adaptive hierarchical sparse grid collocation algorithm for the solution of stochastic differential equations, Journal of Computational Physics 228 (8) (2009) 3084–3113.

[15] E. J. Candès, M. B. Wakin, An introduction to compressive sampling, IEEE signal processing magazine 25 (2) (2008) 21–30.

[16] D. L. Donoho, Compressed sensing, Information Theory, IEEE Transactions on 52 (4) (2006) 1289–1306.

[17] E. J. Candès, J. Romberg, T. Tao, Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information, Information Theory, IEEE Transactions on 52 (2) (2006) 489–509.

[18] A. Doostan, H. Owhadi, A non-adapted sparse approximation of PDEs with stochastic inputs, Journal of Computational Physics 230 (8) (2011) 3015–3034.

[19] L. Mathelin, K. Gallivan, A compressed sensing approach for partial differential equations with random input data, Communications in computational physics 12 (04) (2012) 919–954.

[20] G. Blatman, B. Sudret, Adaptive sparse polynomial chaos expansion based on least angle regression, Journal of Computational Physics 230 (6) (2011) 2345–2367.

[21] X. Yang, H. Lei, N. A. Baker, G. Lin, Enhancing sparsity of Hermite polynomial expansions by iterative rotations, Journal of Computational Physics 307 (2016) 94–109.

[22] J. D. Jakeman, M. S. Eldred, K. Sargsyan, Enhancing ℓ_1 -minimization estimates of polynomial chaos expansions using basis selection, Journal of Computational Physics 289 (2015) 18–34.

[23] N. Alemazkoor, H. Meidani, Divide and conquer: an incremental sparsity promoting compressive sampling approach for polynomial chaos expansions, arXiv preprint arXiv:1606.06611.

[24] H. Rauhut, R. Ward, Sparse Legendre expansions via ℓ_1-minimization, Journal of approximation theory 164 (5) (2012) 517–533.

[25] G. Tang, G. Iaccarino, Subsampled Gauss quadrature nodes for estimating polynomial chaos expansions, SIAM/ASA Journal on Uncertainty Quantification 2 (1) (2014) 423–443.

[26] J. Hampton, A. Doostan, Compressive sampling of polynomial chaos expansions: convergence analysis and sampling strategies, Journal of Computational Physics 280 (2015) 363–386.

[27] D. Xiu, Numerical methods for stochastic computations: a spectral method approach, Princeton University Press, 2010.

[28] Y. Shin, D. Xiu, Nonadaptive quasi-optimal points selection for least squares linear regression, SIAM Journal on Scientific Computing 38 (1) (2016) A385–A411.

[29] M. Lustig, J. M. Santos, J.-H. Lee, D. L. Donoho, J. M. Pauly, Application of compressed sensing for rapid MR imaging, SPARS,(Rennes, France).

[30] J. H. Ender, On compressive sensing applied to radar, Signal Processing 90 (5) (2010) 1402–1414.

[31] J. L. Paredes, G. R. Arce, Z. Wang, Ultra-wideband compressed sensing: channel estimation, Selected Topics in Signal Processing, IEEE Journal of 1 (3) (2007) 383–395.

[32] J. F. Gemmeke, H. Van Hamme, B. Cranen, L. Boves, Compressive sensing for missing data imputation in noise robust speech recognition, Selected Topics in Signal Processing, IEEE Journal of 4 (2) (2010) 272–287.

[33] L. Yan, L. Guo, D. Xiu, Stochastic collocation algorithms using ℓ_1 minimization, International Journal for Uncertainty Quantification 2 (3).

[34] X. Yang, G. E. Karniadakis, Reweighted ℓ_1 minimization method for stochastic elliptic differential equations, Journal of Computational Physics 248 (2013) 87–108.

[35] A. M. Bruckstein, D. L. Donoho, M. Elad, From sparse solutions of systems of equations to sparse modeling of signals and images, SIAM review 51 (1) (2009) 34–81.

[36] S. S. Chen, D. L. Donoho, M. A. Saunders, Atomic decomposition by basis pursuit, SIAM review 43 (1) (2001) 129–159.

[37] E. J. Candès, The restricted isometry property and its implications for compressed sensing, Comptes Rendus Mathematique 346 (9) (2008) 589–592.

[38] M. A. Maleki, D. L. Donoho, R. Gray, A. Montanari, Approximate message passing algorithms for compressed sensing, Stanford University, 2010.

[39] H. Rauhut, Compressive sensing and structured random matrices, Theoretical foundations and numerical methods for sparse recovery 9 (2010) 1–92.

[40] V. Cevher, A. Krause, Greedy dictionary selection for sparse representation, IEEE Journal of Selected Topics in Signal Processing 5 (5) (2011) 979–988.

[41] M. A. Davenport, M. F. Duarte, Y. C. Eldar, G. Kutyniok, Introduction to compressed sensing, Preprint 93 (1) (2011) 2.

[42] D. L. Donoho, M. Elad, Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ_1 minimization, Proceedings of the National Academy of Sciences 100 (5) (2003) 2197–2202.

[43] G. Li, Z. Zhu, D. Yang, L. Chang, H. Bai, On projection matrix optimization for compressive sensing systems, IEEE Transactions on Signal Processing 61 (11) (2013) 2887–2898.

[44] H. Cheng, The fundamentals of compressed sensing, in: Sparse Representation, Modeling and Learning in Visual Recognition, Springer, 2015, pp. 21–53.

[45] D. L. Donoho, M. Elad, V. N. Temlyakov, Stable recovery of sparse overcomplete representations in the presence of noise, Information Theory, IEEE Transactions on 52 (1) (2006) 6–18.

[46] M. Elad, Sparse and redundant representations: From theory to applications in signal and image processing.

[47] M. Elad, Optimized projections for compressed sensing, IEEE Transactions on Signal Processing 55 (12) (2007) 5695–5702.

[48] J. M. Duarte-Carvajalino, G. Sapiro, Learning to sense sparse signals: Simultaneous sensing matrix and sparsifying dictionary optimization, Tech. rep., DTIC Document (2008).

[49] V. Abolghasemi, S. Ferdowsi, B. Makkiabadi, S. Sanei, On optimization of the measurement matrix for compressive sensing, in: Signal Processing Conference, 2010 18th European, IEEE, 2010, pp. 427–431.

[50] L. Zelnik-Manor, K. Rosenblum, Y. C. Eldar, Sensing matrix optimization for block-sparse decoding, IEEE Transactions on Signal Processing 59 (9) (2011) 4300–4312.

[51] S. Tian, X. Fan, Z. Li, T. Pan, Y. Choi, H. Sekiya, Orthogonal-gradient measurement matrix construction algorithm, Chinese Journal of Electronics 25 (1) (2016) 81–87.

[52] E. Van Den Berg, M. Friedlander, SPGL1: A solver for large-scale sparse reconstruction (2007).

[53] J. D. Jakeman, A. Narayan, T. Zhou, A generalized sampling and preconditioning scheme for sparse approximation of polynomial chaos expansions, arXiv preprint arXiv:1602.06879.